arXiv:1703.07028v2 [cs.LO] 27 Mar 2017

Epistemic Modeling with Justifications

Sergei Artemov

The Graduate Center, the City University of New York
365 Fifth Avenue, New York City, NY 10016
sartemov@gc.cuny.edu

March 28, 2017

Abstract

Existing logical models do not fairly represent epistemic situations with fallible justifications, e.g., Russell's Prime Minister example, though such scenarios have long been at the center of epistemic studies. We introduce *justification epistemic models*, *JEM*, which can handle such scenarios. *JEM* makes justifications prime objects and draws a distinction between accepted and knowledge-producing justifications; belief and knowledge become derived notions. Furthermore, Kripke models can be viewed as special cases of *JEM*s with additional assumptions of evidence insensitivity and common knowledge of the model. We argue that *JEM* can be applied to a range of epistemic scenarios in CS, AI, Game Theory, etc.

1 Context and motivations

We make the case that formal epistemology should venture, when needed, beyond modal logic limits. This includes a formal treatment of justifications, which have been at the heart of epistemic discussions since Plato.

Proof systems of justification logic and general purpose classes of models for these systems have been studied in [3, 4, 5, 7, 10, 11, 16, 17, 19] and many other sources. However, for formalizing epistemic scenarios, one needs specific *domain-dependent models*, with additional features that are not necessary for standard soundness and completeness analysis of proof systems.

We introduce *justification epistemic models*, *JEM*, in which justifications are primary objects and a distinction is made between *accepted* and *knowledge-producing* justifications. In *JEM*s, belief and knowledge are derived notions which depend on the status

of supporting justifications. We argue that *JEM*s can work in situations in which standard non-hyperintensional tools (Kripke, topological, algebraic) fail to fairly represent the corresponding epistemic structure.

This paper is self-contained due to streamlining the exposition of justification models, an essential ingredient of *JEM*s.

1.1 Comparisons to other semantics of justifications

Justification logics are obtained from the usual propositional logic by adding new atoms of the sort

$$t{:}F$$

in which t denotes a justification, F a formula, and $t{:}F$ is read informally as

$$t \textit{ is a justification of } F.$$

The first mathematical semantics of justification logic appeared in [2] (cf. also [3]) for the Logic of Proofs LP, namely, a constructive version of the **canonical model** and **arithmetical semantics** in which $t{:}F$ was interpreted as t *is a proof of* F.

In [17], the canonical model construction for LP was formulated in a set-theoretic format, with interpreting proof objects t as sets of formulas t^* and speaking of "$t{:}F$ is true" as $F \in t^*$. This led to generic **basic models** for justification logics (cf. [5, 16] and the current paper). In a basic model $*$,

$$\models_* t{:}F \quad \textit{iff} \quad F \in t^*.$$

To represent multiple possible worlds, basic models were generalized to **modular models** for justification logic J in [5], and then extended to other justification logics in [16]. A modular model is a Kripke model $(W, R, \models)$ with a basic model $*(u)$ at each world $u \in W$. Evaluation of justification assertions $t{:}F$ is defined canonically:

$$u \models t{:}F \quad \textit{iff} \quad F \in t^{*(u)}.$$

Mkrtychev models were also developed in [17] to capture reflexivity in the Logic of Proofs. In Mkrtychev models, propositional and justification evaluations are similar to those in basic models (propositions are Boolean, justification terms are sets of formulas), but the justification assertions $t{:}F$ are evaluated as "t is a justification for F and F is true":

$$\models_* t{:}F \qquad \textit{iff} \qquad F \in t^* \textit{ and } \models_* F.$$

Conceptually, Mkrtychev models are close to reflexive basic models: each reflexive basic model is a Mkrtychev model, each Mkrtychev model contains a reflexive basic model, However, Mkrtychev models are often easier to work with for justification logics with reflection.

Fitting models for LP were introduced and developed in [10] and then extended to other justification logics. Conceptually, the basic model approach reflects only one reason for *not knowing* F: no sufficient justification for F is available. In contrast, Fitting models take into account two reasons for *not knowing* F: the *Kripkean reason*, F fails in some possible world; and the *awareness reason*, no justification for F is available. Technically, a Fitting model is a Kripke model $(W, R, \models)$ with a basic model $*(u)$ at each world $u \in W$ (together with appropriate closure conditions). The difference between Fitting models and modular models lies in the evaluation of justification assertions: for Fitting models,

$$u \models t{:}F \qquad \textit{iff} \qquad F \in t^{*(u)} \textit{ and } R(u) \models F.$$

This secures the *Justification Yields Belief (JYB) property*: $t{:}F$ yields 'F is believed' (i.e., F holds at all accessible worlds).

Modular/basic models have clear motivations and provide an answer to the ontological question *What kind of logical object is a justification?* However, these models might be difficult to work with, and more convenient constructions, such as Mkrtychev and Fitting models, are normally used instead.

Technically, basic models and Mkrtychev models may be regarded as special cases of Fitting models. On the other hand, Fitting models can be identified as modular models with *JYB*, cf. [5]. This provides a natural hierarchy of the aforementioned classes of models:

$$\textit{basic and Mkrtychev models} \subset \textit{Fitting models} \subset \textit{modular models} \subset \textit{JEM}.$$

Even the smallest class, basic models, is already sufficient for mathematical completeness of justification logics. So, the main idea of progressing to Fitting models, modular models, or *JEM*s is not a pursuit of completeness but rather a desire to offer natural and manageable models for a variety of epistemic situations involving evidence, belief, and knowledge.

1.2 Novelties

Several things in this paper might be considered novelties.

1. Separating **accepted justifications** and **knowledge-producing justifications**, and making knowledge and justified belief derived notions. Building an adequate model for Russell's Prime Minister example as a showcase.

2. Streamlined exposition of generic semantics for justifications as sets of formulas, connecting basic justification models with the canonical model of propositional classical logic.

3. Asserting logic J^- instead of J (cf. Section 4.2) as in [4], to be the base system of justification epistemic logic.

2 Preliminaries

Standard modal epistemic models have "propositional" precision, i.e., they do not distinguish sentences with the same truth values at each possible world. The expressive power of such models for analysis of justification, belief, and knowledge is rather limited, and so we have to "go hyperintensional."[1] Specifically, if, at all possible worlds, t is a justification for F

$$\Vdash t{:}F,$$

and G has the same truth value as F

$$\Vdash F \leftrightarrow G,$$

we still cannot conclude that t is a justification for G

$$\nVdash t{:}G.$$

A natural example comes from mathematics: both statements $0 = 0$ and *Fermat's Last Theorem*, FLT, are true (proven) mathematical facts and hence are true at all possible worlds. However, we cannot claim that a proof of $0 = 0$ is a proof of FLT as well.

A sample justification logic analysis of standard epistemic situations (Gettier examples, Red Barn example) is presented in [4] using Fitting models [10] though, due to the relative simplicity of those examples, this analysis could be replicated in a bi-modal language (cf. [20]).

However, we cannot go much farther without adopting a justification framework: the situation changes when we have to represent several conflicting pieces of evidence for a stated fact, cf. the following Russell example of 1912 ([18]):

> *If a man believes that the late Prime Minister's last name began with a 'B,' he believes what is true, since the late Prime Minister was Sir Henry Campbell Bannerman*[2]*. But if he believes that Mr. Balfour was the late Prime Minister, he will still believe that the late Prime Minister's last name began with a 'B,' yet this belief, though true, would not be thought to constitute knowledge.*

To keep it simple, we consider proposition B

$$\textit{the late Prime Minister's last name began with a 'B,'}$$

with two justifications for B: the right one r and the wrong one w; the agent chooses w as a reason to believe that B holds.

[1] From [9]: "Hyperintensional contexts are simply contexts which do not respect logical equivalence."
[2] which was true in 1912

There is a mathematical version of the story with a true proposition and its two justifications; one is correct, the other is not. Consider the picture[3]:

$$\frac{1\not{6}}{\not{6}4} = \frac{1}{4}. \tag{1}$$

The true proposition is "16/64 = 1/4," the right justification is dividing both the numerator and the denominator by 16, and the wrong (but shorter and easier) justification is simplifying as in (1).

To formalize Russell's scenario in modal logic (cf. [20]), we introduce two modalities: **K** for knowledge and **J** for justified belief. In the real world,

- B holds;
- $\mathbf{J}B$ holds, since the agent has a justification w for B;
- $\mathbf{K}B$ does not hold;

thus yielding the set of assumptions

$$\Gamma = \{B,\ \mathbf{J}B,\ \neg\mathbf{K}B\}.$$

However, Γ doesn't do justice to Russell's scenario: the right justification r is not represented and Γ rather corresponds to the same scenario but lacking r. The intrinsic epistemic structure of the example is not respected.

Within the *JEM* framework, introduced in this paper, we provide a model for Russell's Prime Minister example which, we wish to think, fairly represents its intrinsic epistemic structure.

3 Generic logical semantics of justifications

What kinds of logical objects are justifications? When asked in a mathematical context "what is a predicate?," we have a ready answer: a subset of a Cartesian product of the domain set. Within an exact mathematical theory, there should be a similar kind of answer to the question "what is a justification?".

We consider this question in its full generality which, surprisingly, yields a clean and meaningful answer. We assume the language of justification logic consists of two disjoint sets of syntactic objects:

1. a set of **justification terms** Tm;

[3]which the author saw on the door of the Mathematics Support Center at Cornell this February.

2. a set of **formulas** Fm, built inductively from propositional atoms using Boolean connectives and the justification formula formation rule: if F is a formula, $F \in Fm$, and t a justification term, $t \in Tm$, then $t{:}F$ is again a formula, $t{:}F \in Fm$.

The meaning assigned to formulas is a classical truth value, 0 for *false* and 1 for *true*, and we retain classical logic behavior for propositional connectives. The key item is to give meaning to justification terms, and this will be a *set of formulas* interpreted as *the set of formulas for which it is a justification*. Here is a formal definition.

Definition 1 [Basic Model] A *basic model*, simply called $*$, consists of an interpretation of the members of Fm, and an interpretation of the members of Tm.

The interpretation of a formula in a basic model is a truth value. That is,

$$* : Fm \mapsto \{0, 1\}.$$

We assume $(X \to Y)^* = 1$ if and only if $X^* = 0$ or $Y^* = 1$, and similarly for the other Boolean connectives. Let also $\models_* X$ stand for $X^* = 1$.

We interpret justification terms as *sets of formulas*. That is,

$$* : Tm \mapsto 2^{Fm}.$$

Our final requirement connects the two mapping roles that $*$ plays in a basic model. For any $X \in Fm$ and any $t \in Tm$,

$$\models_* t{:}X \quad \text{if and only if} \quad X \in t^*.$$

It is easy to check that any mapping $*$ from propositional letters to truth values, and from justification terms to sets of formulas determines a unique basic model.

So far, a basic model is nothing but a classical propositional model in which justification assertions $t{:}F$ are treated as *independent propositional atoms*.

Note that while propositions are interpreted semantically, as truth values, justifications are interpreted syntactically, as sets of formulas. This is a principal *hyperintensional* feature: a basic model may treat distinct formulas F and G as equal, i.e. $F^* = G^*$, but still be able to distinguish justification assertions $t{:}F$ and $t{:}G$, e.g., when $F \in t^*$, but $G \notin t^*$ yielding $\models_* t{:}F$ but $\not\models_* t{:}G$.

Definition 2 Let $S \subseteq Fm$ and $X \in Fm$. We write $S \vdash X$ if formula X is derivable from the set S of formulas in classical logic treating justification assertions $t{:}F$ as propositional atoms, and with *Modus Ponens* as the only rule of inference. We say that S is *consistent* if $S \nvdash \bot$.

A basic model of S is merely a possible world containing S in the canonical model, i.e., a maximal consistent set Γ of formulas, with the convenience agreement to read $t{:}F \in \Gamma$ as $F \in \{X \mid t{:}X \in \Gamma\}$. In this respect, basic models and the canonical model

are slightly different but obviously equivalent ways of presenting the same object. When we move to more sophisticated models (Fitting models, modular models), the advantage of dealing with sets and operations (e.g. basic models) over logical conditions (e.g. the canonical model) becomes clear.

Definition 3 For $S \subseteq Fm$, $BM(S)$ is the *class of all basic models of* S.

Theorem 1 (Generic Soundness and Completeness) *Each set of formulas S is sound and complete with respect to its class of basic models $BM(S)$. In other words, $S \vdash F$ if and only if F is true in each basic model of S.*

Proof. This theorem is merely a reformulation of the soundness and completeness of classical propositional logic with hypotheses. Indeed, if $S \vdash F$ and $\models_* S$, then $\models_* F$ since propositional derivations respect validity.

If $S \nvdash F$, then there is a Boolean evaluation $*$ which makes all formulas from S true, $S^* = 1$, and F false, $F^* = 0$. In this case there are two types of atomic propositions: propositional letters P and justification assertions $t{:}X$. Define

$$t^* = \{X \mid (t{:}X)^* = 1\}$$

and note that $(t{:}X)^* = 1$ iff $X \in t^*$. Therefore, $*$ is a propositional evaluation and $*$ is a basic model yielding the same truth values of atomic formulas P and $t{:}X$. Since $S^* = 1$ and $F^* = 0$, we have $\models_* S$ and $\not\models_* F$ for basic model $*$. □

An easy corollary: $\vdash F$ iff F is a tautology (with $t{:}X$'s as distinct propositional atoms).

Example 1 We work with classical logic over Fm, as in Definition 2, taking $S = \emptyset$.

1. For any justification term t,

   $$\nvdash t{:}F.$$

   Straightforward, since $t{:}F$ is not a propositional tautology. For a specific counter-model put $t^* = \emptyset$ for each term $t \in Tm$, which makes $\not\models_* t{:}F$.

2. For any propositional letter P, and term t,

   $$\nvdash t{:}P \to P.$$

   Likewise, this holds because $t{:}P \to P$ is not a propositional tautology. Specifically, put $t^* = Fm$ and $P^* = 0$, with other assignments being arbitrary. In this model, all justification assertions are true, but $t{:}P \to P$ is false.

3. For any propositional letter P, and term t,

$$\nvdash P \to t{:}P.$$

Again, this holds since $P \to t{:}P$ is not a propositional tautology. For example, put $t^* = \emptyset$ and $P^* = 1$. In this model, t is not a justification for P (i.e., $\not\models_* t{:}P$) and $P \to t{:}P$ is false.

4. A somewhat less trivial example illustrating hyperintensionality: for a justification variable x and formula F

$$\nvdash x{:}F \to x{:}(F \wedge F).$$

The high-level argument is the same: formulas $x{:}F$ and $x{:}(F \wedge F)$ from Boolean evaluation point of view can be regarded as distinct propositional variables. Hence $x{:}F \to x{:}(F \wedge F)$ is not a tautology. For a countermodel, take $x^* = \{F\}$. Then $\models_* x{:}F$, but $\not\models_* x{:}(F \wedge F)$. This demonstrates hyperintensionality of justification logic base, since F and $F \wedge F$ are provably equivalent, but not $x{:}F$ and $x{:}(F \wedge F)$.

4 Basic Justification Logic J^-

Within the Justification Logic framework, there are two sorts of logical objects: justification terms Tm and formulas Fm. Let us become more specific about both.

- For Tm, reserve a set of justification constants $a, b, c, \ldots$ with indices, and variables $x, y, z, \ldots$ with indices. Justification terms are built from constants and variables by a binary operation $\cdot$ (application).
- Formulas are built from propositional letters $p, q, r, \ldots$ (with indices) and Boolean constant $\bot$ (falsum) by the standard Boolean connectives $\wedge, \vee, \to, \neg$ with a new formation rule: *whenever t is a justification term and F is a formula, $t{:}F$ is a formula (with the informal reading "t is a justification for F").*

The *logical system* J^- consists of two groups of postulates.

- **Background logic**: axioms of classical propositional logic, rule *Modus Ponens*.
- **Application**: $s{:}(F \to G) \to (t{:}F \to [s \cdot t]{:}G)$.

Basic models corresponding to J^- are the ones in which the application axiom holds. They can be specified by a natural combinatorial condition.

Definition 4 For sets of formulas S and T, we define

$$S \triangleright T = \{F \mid G \to F \in S \text{ and } G \in T \text{ for some } G\}.$$

Informally, $S \triangleright T$ is the result of applying *Modus Ponens* once to all members of S and of T (in a given order).

Theorem 2 *$BM(\mathsf{J}^-)$ is the class of basic models with the following closure condition*

$$s^* \triangleright t^* \subseteq (s{\cdot}t)^*. \tag{2}$$

Proof. Let us assume the closure condition (2) and check the validity of the application axiom. Indeed, $\models_* s{:}(F \to G)$ and $\models_* t{:}F$ yield $(F \to G) \in s^*$ and $F \in t^*$. By the closure condition, $G \in [s{\cdot}t]^*$, i.e., $\models_* [s{\cdot}t]{:}G$.

Now assume the application axiom and derive the closure condition (2). Let $(F \to G) \in s^*$ and $F \in t^*$. By definitions, this yields $\models_* s : (F \to G)$ and $\models_* t : F$. By the application axiom, $\models_* [s{\cdot}t]{:}G$, hence $G \in [s{\cdot}t]^*$. □

Example 2 None of the formulas from Example 1: $t{:}F$, $t{:}P \to P$, $P \to t{:}P$, $x{:}F \to x{:}(F \wedge F)$ is derivable in J^-. Indeed, every specific evaluation from Example 1.1–3 satisfies the closure condition (2), hence their countermodels are J^--models. Consider the latter formula (4). Put $x^* = \{F\}$ and $t^* = Fm$ for all other terms t. The closure condition (2) holds vacuously, hence $*$ is a J^--model. Obviously, $\models_* x{:}F$ and $\not\models_* x{:}(F \wedge F)$.

Constants in justification logic are used to denote justifications of assumptions, in particular, axioms. Indeed, as we have already seen in Example 2, no formula $t : F$ is derivable in J^-. In particular, no logical axiom is assumed justified in J^- which is not realistic.

Definition 5 A set S of formulas is *reflexive* if for each $s{:}t{:}F \in S$, $t{:}F$ is also in S. By constant specification CS we understand a reflexive set of formulas of the type

$$c_n{:}c_{n-1}{:}c_{n-2}{:}\ \ldots\ c_1{:}A$$

where A is a J^--axiom and c_i are justification constants. The major classes of constant specifications:

- *empty*,
- *total*— (each constant is a justification for each axiom),
- *axiomatically appropriate* (each axiom has a justification at any depth),
- *injective* (each constant justifies at most one axiom).

Let CS be a constant specification. Then by $\mathsf{J}^-(CS)$, we understand J^- with additional axioms CS. A *CS-model* is a model in which all formulas from CS hold.

Corollary 1 *Basic models for $\mathsf{J}^-(CS)$ are the basic CS-models for J^-. $\mathsf{J}^-(CS)$ is sound and complete with respect to the class of its basic models.*

4.1 Sharp models

In closure condition (2) from Theorem 2, one cannot, generally speaking, replace the inclusion "$\subseteq$" by the equality "$=$" without violating completeness Theorem 1.

Indeed, fix a justification constant 0 and consider logic

$$\mathcal{L} = \mathsf{J}^- + \{\neg 0{:}F \mid F \in Fm\}.$$

Informally, justification 0 receives empty evaluation in any basic model, $0^* = \emptyset$. We claim that formula $F = \neg[0{\cdot}0]{:}P$ is not derivable in $\mathcal{L}$, but is true in any basic model of $\mathcal{L}$ with the closure condition $s^* \triangleright t^* = (s{\cdot}t)^*$. To show that $\mathcal{L} \not\vdash F$, it suffices to find a basic model for $\mathcal{L}$ in which F is false. Consider a basic model $\sharp$ such that $0^\sharp = \emptyset$ and $t^\sharp = Fm$ for any other justification term t. Obviously, the closure condition from Theorem 2, together with $0^\sharp = \emptyset$, is met. Therefore, $\sharp$ is a basic model of $\mathcal{L}$. It is immediate that F is false in $\sharp$, since $[0{\cdot}0]^\sharp = Fm$. On the other hand, for any basic model of $\mathcal{L}$ with the closure condition $(0{\cdot}0)^* = 0^* \triangleright 0^*$, $(0{\cdot}0)^* = \emptyset$ since $0^* = \emptyset$ and $\emptyset \triangleright \emptyset = \emptyset$.

Definition 6 *Sharp* basic models are those in which the application closure condition has the form

$$(s{\cdot}t)^* = s^* \triangleright t^*. \tag{3}$$

One can prove that J^- and all its constant specification extensions of $\mathsf{J}^-(CS)$ are complete with respect to their sharp basic models. We skip the proof as not essential for the purposes of this paper.

Note that a sharp model is completely defined by evaluations of atomic propositions and atomic justifications.

4.2 Other justification logics

There is a whole family of justification logics and they all extend J^-; the reader is referred to [4, 7] for details. Here we list just the main systems of justification logic for a general orientation.

Logic J is obtained from J^- by adding a new operation on justifications '+' and the Sum principle

$$s{:}F \vee t{:}F \to [s + t]{:}F.$$

Logics JD, JT, $\mathsf{J4}$, $\mathsf{J5}$, etc., are obtained by adding the corresponding combination of principles

$$D = \neg t{:}\bot,$$

$$T = t{:}F \to F,$$

$$4 = t{:}F \to !t{:}t{:}F,$$

$$5 = \neg t{:}F \to ?t{:}\neg t{:}F.$$

Historically, the first justification logic was the Logic of Proofs LP [2, 3] which in these notations can be identified as JT4. By now, the family of justification logics has grown to be infinite, cf. [12].

5 Justification Epistemic Models

We need more expressive power to capture epistemic differences between justifications and their use by the knower. Some justifications are knowledge-producing, some are not. The agent makes choices on which justifications to base an agent's beliefs/knowledge and which justifications to ignore in this respect. These actions are present in epistemic scenarios, out of which we will primarily focus on Russell's example, which has it all:

- there are justifications w (Balfour was the late prime minister) and r (Bannerman was the late prime minister) for B;
- r is knowledge-producing whereas w is not;
- the agent opts to base his belief on w and ignores r;
- the resulting belief is evidence-based, but is not knowledge.

Fix $\mathsf{J}^-(CS)$ for some axiomatically appropriate constant specification CS.

Definition 7 A set X of justification terms is *properly closed* if X contains all constants and is closed under applications. If Y is a set of justification terms, then by $\overline{Y}$ we mean the proper closure of Y, i.e., the minimal properly closed superset of Y.

Definition 8 A (basic) *Justification Epistemic Model* (JEM) is

$$(*, \mathcal{A}, \mathcal{E})$$

where

- $*$ is a basic $\mathsf{J}^-(CS)$-model;
- $\mathcal{A} \subseteq Tm$ is a properly closed set $\mathcal{A}$ of *accepted justifications*[4];
- $\mathcal{E} \subseteq Tm$ is a properly closed sets $\mathcal{E}$ of *knowledge-producing justifications*; all knowledge-producing justification should be factive, $\models_* t{:}F \to F$ for each F and each $t \in \mathcal{E}$.

Both sets $\mathcal{A}$ and $\mathcal{E}$ contain all constants. This definition presumes that constants in a model are knowledge-producing and accepted.

[4]Consistency $\models_* \neg t{:}\bot$ for any $t \in \mathcal{A}$ may or may not be required.

Definition 9 Sentence F is *believed* if there is $t \in \mathcal{A}$ such that $\models_* t{:}F$. Sentence F is *known* if there is $t \in \mathcal{A} \cap \mathcal{E}$ such that $\models_* t{:}F$.

By *ground term* we understand a term containing no (justification) variables. In other words, a term is ground iff it is built from justification constants only.

Sets of accepted and knowledge-producing justifications overlap on ground terms but otherwise are in a general position[5]. There may be accepted but not knowledge-producing justifications, and vice versa. So, *JEM*s do not analyze **why** certain justifications are knowledge-producing (knowledge-grade) or accepted, but rather provide a formal framework that accommodates these notions.

5.1 Injective justifications

The notions of *accepted* and *knowledge-producing* justifications should be utilized with some caution. Imagine a justification t for F (i.e., $t : F$ holds) and for G ($t : G$) such that, intuitively, t is a knowledge-producing justification for F but not for G. Is such a t knowledge-producing, trustworthy, acceptable for a reasonable agent? The answers to these questions seem to depend on F and G, and if we prefer to handle justifications as objects rather than as justification assertions, it is technically convenient to assume that justifications are *injective*, or statement-specific:

there is at most one formula F such that $t{:}F$ holds.

Conceptually, by going injective with justifications, one does not lose generality: if p is a proof of F and of something else, then the same p with a designated statement F, symbolically, a pair (p, F), can be regarded as an injective/single conclusion proof of F. So, with injective justifications without losing generality we gain the ability to go back and forth between justifications and statements.

In model $\mathcal{R}$ for the Russell Example, Section 6, all justifications are injective.

Note that J^- is not complete with respect to the class of basic models which are both sharp and injective (as model $\mathcal{R}$ for the Russell Example). Indeed, consider formula F,

$$F = \neg(x{:}(P \to Q) \land y{:}P \land [x \cdot y]{:}R)$$

where P, Q, R are distinct propositional letters and x, y justification variables. Obviously, F holds in any basic model $*$ which is sharp and injective. Imagine a sharp injective $*$ in which $x{:}(P \to Q)$ and $y{:}P$ hold. In such $*$, $[x \cdot y]^* = \{Q\}$, hence both $\neg[x \cdot y]{:}R$, and F hold. On the other hand, F is not derivable in J^-, e.g., F fails in the basic model $*$ with $x^* = \{P \to Q\}$, $y^* = \{P\}$, and $t^* = Fm$ for any other t (check closure condition (2)!).

[5]In principle, one could consider smaller sets $\mathcal{A}$, which would correspond to the high level of skepticism of an agent who does not necessarily accept logical truths (axioms) as justified. We leave this possibility for further studies.

So, "sharp and injective" justification tautologies constitute a proper extension of J^-. It would be interesting to find its complete axiomatization. It would seem that functionality axioms from [8, 14, 15] might be relevant here.

6 Russell scenario as a *JEM*

Consider the version of J^- in language with two justification variables w and r, one propositional letter B, and injective constant specification CS:

$$c_n{:}A \in CS \text{ iff } A \textit{ is an axiom and } n \textit{ is the Gödel number of } A.$$

Define a model $*$ such that

- $B^* = 1$, i.e., $\models_* B$;
- $c_n^* = \{A\}$ if A is an axiom and n is the Gödel number $|A|$ of A, and $c_n^* = \emptyset$ otherwise;
- $w^* = r^* = \{B\}$, e.g., $\models_* r{:}B$ and $\not\models_* r{:}F$ for any F other than B (the same for w);
- Application is sharp: $(s{\cdot}t)^* = s^* \triangleright t^*$.

A *JEM* $\mathcal{R}$ (for Russell's scenario) is $(*, \mathcal{A}, \mathcal{E})$ with

- $\mathcal{A} = \overline{\{w\}}$, i.e., the set of accepted justifications is $\{w\}$, properly closed;
- $\mathcal{E} = \overline{\{r\}}$, i.e., the set of knowledge-producing justifications is $\{r\}$, properly closed.

Though the idea behind $\mathcal{R}$ is quite intuitive, we need to fill in some technical details: extending truth evaluations to all terms and formulas and checking closure conditions.

6.1 Technicalities of the model

Put $c^*_{|A|} = \{A\}$ for each axiom A of $\mathsf{J}^-(CS)$. Technically, this is an inductive definition. Base: $c_0^* = \emptyset$, given 0 is not a Gödel number of any formula. Inductive step: suppose n is the Gödel number $|F|$ of some formula F. If F is an axiom of J^-, put $c_n^* = \{F\}$. If $F = c_k{:}G$ for some c_k and G, then, by monotonicity of Gödel numbering, $k < n$, hence c_k^* is defined. If $c_k^* = \{G\}$, then $c_k{:}G$ is an axiom of $\mathsf{J}^-(CS)$ and we can put $c_n^* = \{F\}$. In all other cases, $c_n^* = \emptyset$.

Since application is sharp, the evaluation of each term is, at most, a singleton. Together with Boolean truth tables, this determines the truth value of any formula.

Lemma 1 *Each $t \in Tm$ is factive, $\models_* t{:}F \to F$.*

Proof. Induction on t. Assume $\models_* t{:}F$; that means $t^* = \{F\}$. If t is w or r, then F is B, which is true in the model $*$. If t is a constant, then F is an axiom and hence true in $*$. The induction step corresponds to application, which preserves the truth of justified formulas. □

From Lemma 1 it follows that all accepted justifications are consistent and all knowledge-producing justifications are factive. Therefore, $\mathcal{R} = (*, \mathcal{A}, \mathcal{E})$ is indeed a *JEM*.

6.2 Main result

Theorem 3 *In model $\mathcal{R}$, sentence B is true, justified and believed, but not known.*

Proof. In model $\mathcal{R}$, sentence B is

- true, since $\models_* B$;
- justified, since $\models_* w{:}B$;
- believed, since $w \in \mathcal{A}$.

We have to show that for any justification $g \in \mathcal{A} \cap \mathcal{E}$, $\not\models_* g{:}B$.

Consider an auxiliary basic model $\bullet$ which is the same as $*$ but with $B^\bullet = 0$, i.e., the truth value of B is flipped from 'true' to 'false.' In particular, application in $\bullet$ is sharp.

Lemma 2 *For each justification term t,*

$$t^* = t^\bullet.$$

Proof. The inductive process (based on sharp application) of evaluating all justifications given evaluations of atomic justifications operates only with formulas of type $t{:}F$ and starts with the same initial set of such formulas in $*$ and $\bullet$. Hence the results of these processes in $*$ and $\bullet$ coincide.

□

In particular, for all $g \in \mathcal{A} \cap \mathcal{E}$, $g^* = g^\bullet$, and if $\models_* g{:}B$, then $\models_\bullet g{:}B$ as well.

Lemma 3 *Each $g \in \mathcal{A} \cap \mathcal{E}$ is factive in $\bullet$, i.e., $\models_\bullet g{:}F \to F$.*

Proof. All $g \in \mathcal{A} \cap \mathcal{E}$ are obtained from constants by application. By construction, if $\models_\bullet c{:}X$, then $X \in c^*$ and X is an axiom, hence true. Application obviously preserves factivity.

□

To complete the proof of Theorem 3, suppose $\mathcal{R} \models g{:}B$, i.e., $\models_* g{:}B$, for some $g \in \mathcal{A} \cap \mathcal{E}$. By Lemma 2, $\models_\bullet g{:}B$, and, by Lemma 1, $\models_\bullet B$, which is not the case.

□

6.3 Can Russell's scenario be made modal?

One could try to express Russell's scenario in a modal language by introducing the justified belief modality

$$\mathbf{J}F \quad \Leftrightarrow \quad \textit{there is } t \in \mathcal{A} \textit{ such that } \models t{:}F,$$

and the knowledge-producing modality

$$\mathbf{E}F \quad \Leftrightarrow \quad \textit{there is } t \in \mathcal{E} \textit{ such that } \models t{:}F,$$

and by stipulating that F is known iff F is both accepted and supported by a knowledge-producing justification:

$$\mathbf{K}F \quad \Leftrightarrow \quad \mathbf{J}F \wedge \mathbf{E}F.$$

This, however, fails, since both $\mathbf{J}B$ and $\mathbf{E}B$ hold in $\mathcal{R}$, but $\mathbf{K}B$ does not. We are facing a Gettier-style phenomenon, when a proposition is supported by a knowledge-producing justification (hence true), and believed, but not known (since knowledge-producing and accepted justifications for B are different). This once again illustrates the limitations of modal language in tracking and sorting justifications.

7 Multi-world *JEM*s and epistemic Kripke models

A multi-world version of *JEM* is

$$\mathcal{M} = (W, *, \mathcal{A}, \mathcal{E})$$

such that W is a nonempty set of states (possible worlds) and for each $u \in W$, the corresponding $(*, \mathcal{A}, \mathcal{E})$ is a single-state *JEM*.

Let us outline what simplifying assumptions would make a *JEM* $\mathcal{M} = (*, \mathcal{A}, \mathcal{E})$ a Kripke model for modal logic.

First, we have to convert justification assertions $t{:}F$ into modal epistemic assertions $\Box F$. A natural way of accomplishing this would be postulating a *justification indifference* property:

$$t{:}F \;\leftrightarrow\; s{:}F$$

holds at each world.

Second, we have to adopt an additional *fully explanatory property* of the model. We define the accessibility relation R on W as

$$uRv \quad \textit{iff} \quad \textit{for all } t \textit{ and } F, u \models t{:}F \Rightarrow v \models F.$$

This standard definition secures that

$$u \models t{:}F \quad \Rightarrow \quad R(u) \models F,$$

but not "$\Leftrightarrow$," which is a hidden assumption in Kripke models, cf. [6]. We have to make this assumption explicit. A model $\mathcal{M}$ is *fully explanatory*, if

$$R(u) \models F \quad \Rightarrow \quad \textit{for some } t, u \models t{:}F,$$

cf. [10] in which this property was introduced and studied. The fully explanatory property can be viewed as a propositional form of knowledge of the model, cf. [6].

It is now straightforward that modal logic T corresponds to *JEM*s with justification indifference and the fully explanatory property (a narrow special case of *JEM*s). To obtain more conventional $\mathsf{S5}$-models, we have to adopt yet another additional assumption of *decidability of justifications*, but leave this to further studies.

8 Discussion

Model $\mathcal{R}$ fairly represents a situation with two justifications, right and wrong, for a given proposition and, in this respect, serves as a showcase for modeling epistemic situations with fallible justifications.

There are many natural open questions that indicate possible research directions:

- Are justification assertions checkable, decidable for an agent?
- Is the property of a justification to be knowledge-producing checkable by the agent?
- In multi-agent cases, how much do agents know about each other and about the model?
- Do agents know each other's accepted and knowledge-producing justifications?
- What is a complete axiomatization of sharp and injective models over J^-?
- What is the complexity of these new justification logics and what are their feasible fragments which make sense for epistemic modeling?

9 Acknowledgements

The author is grateful to Melvin Fitting, Vladimir Krupski, Elena Nogina, and Tudor Protopopescu for helpful suggestions. Special thanks to Karen Kletter for editing and proofreading this text.

References

[1] S. Artemov. Logic of proofs. *Annals of Pure and Applied Logic*, 67(1–3): 29-59, 1994.

[2] S. Artemov. *Operational modal logic*, Technical Report MSI 95-29. Cornell University, 1995.

[3] S. Artemov. Explicit provability and constructive semantics. *Bulletin of Symbolic Logic*, 7(1):1–36, 2001.

[4] S. Artemov. The logic of justification. *The Review of Symbolic Logic*, 1(4):477–513, 2008.

[5] S. Artemov. The Ontology of Justifications in the Logical Setting. *Studia Logica* 100(1–2):17–30. 2012.

[6] S. Artemov. Knowing the Model. *ArXiv Preprint*,arXiv:1610.04955, 2016.

[7] S. Artemov and M. Fitting. Justification Logic. *The Stanford Encyclopedia of Philosophy (Winter 2015 Edition)*, Edward N. Zalta (ed.),
URL = http://plato.stanford.edu/archives/win2015/entries/logic-justification, 2015.

[8] S. Artemov and T. Strassen. *Functionality in the basic logic of proofs.* Technical Report IAM 93-004, Department of Computer Science, University of Bern, Switzerland, 1993.

[9] Max J. Cresswell. Hyperintensional logic. *Studia Logica* 34(1):25–38, 1975.

[10] M. Fitting. The logic of proofs, semantically. *Annals of Pure and Applied Logic*, 132(1):1–25, 2005.

[11] M. Fitting. Possible world semantics for first-order logic of proofs. *Annals of Pure and Applied Logic*, 165(1):225–240, 2014.

[12] M. Fitting. Modal logics, justification logics, and realization. *Annals of Pure and Applied Logic*, 167(8):615-648, 2016.

[13] E. Gettier. Is justified true belief knowledge? *Analysis*, 23:121–123, 1963.

[14] V. Krupski. Operational logic of proofs with functionality condition on proof predicate. In *International Symposium on Logical Foundations of Computer Science*, pp. 167–177. Springer Berlin Heidelberg, 1997.

[15] V. Krupski. The single-conclusion proof logic and inference rules specification. *Annals of Pure and Applied Logic*, 113(1):181–206, 2002.

[16] R. Kuznets and T. Struder. Justifications, Ontology, and Conservativity. In *Advances in Modal Logic*, Volume 9, edited by T. Bolander, T. Braner, S. Ghilardi, and L. Moss, 437–458. College Publications, 2012.

[17] A. Mkrtychev. Models for the Logic of Proofs. In S. Adian and A. Nerode (eds.), *Logical Foundations of Computer Science, LFCS'97*, Proceedings (Lecture Notes in Computer Science: Volume 1234), Berlin: Springer, pp. 266–275, 1997.

[18] B. Russell. *The Problems of Philosophy*. London: Williams and Norgate; New York, NY: Henry Holt and Company, 1912.

[19] I. Sedlár. Justifications, Awareness and Epistemic Dynamics. In In S. Artemov and A. Nerode (eds.), *Logical Foundations of Computer Science, LFCS 2013*, Proceedings (Lecture Notes in Computer Science: Volume 7734), Berlin/Heidelberg: Springer, pp. 307–318, 2013.

[20] T. Williamson. A note on Gettier cases in epistemic logic. *Philosophical Studies* 172(1):129–140, 2015.

[21] T. Yavorskaya. Logic of proofs and provability. *Annals of Pure and Applied Logic*, 113(1): 345–372, 2001.